arXiv:1702.05413v1 [cs.CV] 17 Feb 2017

3D Cell Nuclei Segmentation with Balanced Graph Partitioning

Julian Arz[1], Peter Sanders[1], Johannes Stegmaier[2], and Ralf Mikut[2]

[1]Institute for Theoretical Informatics, KIT, Karlsruhe, Germany.
[2]Institute for Applied Computer Science, KIT, Karlsruhe, Germany.

Abstract

Cell nuclei segmentation is one of the most important tasks in the analysis of biomedical images. With ever-growing sizes and amounts of three-dimensional images to be processed, there is a need for better and faster segmentation methods. Graph-based image segmentation has seen a rise in popularity in recent years, but is seen as very costly with regard to computational demand. We propose a new segmentation algorithm which overcomes these limitations. Our method uses recursive balanced graph partitioning to segment foreground components of a fast and efficient binarization. We construct a model for the cell nuclei to guide the partitioning process. Our algorithm is compared to other state-of-the-art segmentation algorithms in an experimental evaluation on two sets of realistically simulated inputs. Our method is faster, has similar or better quality and an acceptable memory overhead.

1 Introduction

The ultimate objective of developmental biology is to understand the development of a single symmetric cell to a complex three-dimensional organism. A big step towards this goal are recent advances in light microscopy technology, which have paved the way for the large-scale *in vivo* investigation of biological structure. For an example, the experiments of Tomer et al. [49] on the embryonic development of *Drosophilia* result in more than one million images with a total size of about 10 TB per embryo.

Already the amounts and sizes of these images make it difficult for them to be manually analyzed. Another challenge is posed by the imaging quality. The images suffer from the presence of noise, low contrast and a blurring effect induced by the physical properties of the optical system. When combining the image stacks to three-dimensional images, a degradation of image quality with increasing axial depth can be observed. This phenomenon is due to the scattering of light in living tissue and shadowing of tightly clustered objects. Another obstacle which has to be overcome is the anisotropic resolution. The resolution[1] in axial direction is typically lower than in lateral direction, by a factor of three to five. This renders it difficult to reconstruct a three-dimensional boundary of an object.

To overcome these challenges, automated methods have to be developed which process and analyze these images. The automation also has the advantage that it frees the image analysis

[1]i.e. the number of voxels per unit of length

process from the researcher's subjective bias and guarantees the reproducibility of the results. Typically, image processing methods consist of a pipeline of multiple tools: from multiview image restoration to object segmentation, annotation and tracking to the modeling of entire processes. Each stage relies on the results of the previous ones. It is important that each tool have results of high quality, despite the aforementioned image deficiencies. Furthermore, it is natural to demand that the total time to analyze the images is in the order it takes for them to be taken, such that computation time does not evolve into a bottleneck of developmental biology research. Therefore, new algorithms have to be efficient, both in quality and in time.

One of the most important steps in nuclei image analysis is *segmentation*. The goal of segmentation is to accurately delineate the boundaries of all observed cell nuclei. It translates the pixel-based representation of the data to an object-based one [28] and is therefore crucial for downstream tasks such as cell state annotation, cell tracking or lineaging.

Many multi-purpose image segmentation algorithms are based on a graph representation of the image [42, 3]. By mapping each pixel to a node in a graph, the image segmentation problem can be reduced to graph theoretic algorithms. The drawback of the employment of the graph-based techniques is their high memory overhead, which, in the context of cell segmentation, makes them suffer most from the steady increase in the size and the amount of images to be processed.

In our work, we develop a new graph-based 3D cell nuclei segmentation algorithm which overcomes the current limitations of these methods. We extract foreground regions from the background with a fast and effective binarization. We perform graph-based segmentation successively on each connected foreground component. This approach solves the problem of the graph representation induced memory overhead, because the size of these components does not directly depend on the image size. The splitting phase uses a probabilistic model based on domain knowledge to make decisions. We perform an experimental evaluation on two sets of realistically simulated inputs. The experiments confirm that our algorithm is both efficient and allows a high quality segmentation even of tightly clustered cells.

Confocal Fluorescence Microscopy. Biological structures at cell level are mostly transparent. Therefore, fluorescent markers, such as the green fluorescent protein (GFP), are introduced into the structures to be observed. When illuminated with an excitation wavelength, the fluorophore emits light of a longer wavelength. In conventional fluorescent microscopes, the entire specimen is illuminated. This poses a problem because emitted light from out-of-focus areas reduces the signal contrast from the in-focus plane [24].

Confocal laser-scanning microscopy (CLSM) [34] limits the field of illumination to the region in focus. A point of light scans the specimen by moving over all planes, lines and points. This allows for optical sectioning of an organism, i.e. creating a 3D image. Confocal fluorescence microscopy was among the most frequently used techniques in the last decades. However, this technique still suffers from the limited accessible depths, because the laser has to penetrate the entire specimen. Another problem is the effect of photobleaching, a process in which the markers are left unable to fluoresce. These limitations hinder its service for long-time detailed imaging of developing organisms.

In contrast, in selective plane illumination microscopy (SPIM) [21], the entire in-focus plane of the sample is illuminated orthogonally to the detection axis. This technique increases the feasible depth of the sample, as the illumination light does not have to penetrate the sample up to the in-focus plane. It also increases the attainable axial resolution. Furthermore, only the

current plane is affected by photobleaching. A further reduction in phototoxicity is achieved with digital scanned laser light sheet fluorescence microscopy (DSLM) [27]. These advances make it possible to image organisms over long periods of time, with high spatiotemporal resolutions, enabling new studies of cell behavior in complex developing specimens [28]. For an example, Keller et al. [27] record entire zebrafish and analyze cell nuclei positions and movement over the first 24 hours of embryogenesis. The study produces stacks of about 400 images with a size of 8 MP and a temporal resolution of up to 60 s. Tomer et al. [49] combine light-sheet microscopy with multi-view imaging. Their experiments on the embryonic development of *Drosophilia* result in more than one million images with a total size of about 10 TB per embryo.

1.1 Related Work

Due to the difficulties of the cell nuclei segmentation task, most algorithms apply a sequence of processing methods. We explain some of these methods, their goals and their relation to each other.

1.1.1 Graph-based Image Segmentation

A straightforward approach for graph-based image segmentation is to define edges with weights based on the similarity between neighboring pixels and then minimize the cut (maximize dissimilarity) between two regions. However, this approach is biased towards unnaturally small regions, because these have less outgoing edges and hence often a smaller cut. Therefore, a second important graph-based image segmentation technique is the minimization of the *normalized cut* [42]. This criterion maximizes both the similarity inside a region and the dissimilarity between regions.

Another popular method is the graph cut framework (see Section 1.2.4). Al-Kofahi et al. [1] apply it both for image binarization and splitting of nuclei clusters. In recent experimental evaluations, graph cut-based cell nuclei segmentation methods have ranked best regarding the segmentation quality [45].

The graph cut multi-label techniques' running times grow linearly with the number of labels, i.e. the number of image objects. This drawback makes additional work necessary to render them practicable for more than 20 labels [1].

1.1.2 Binarization

Often a first step in image analysis is a binarization of the image, with the aim to distinguish the foreground from the background.

Image binarization algorithms such as Otsu's method [36] or adaptive iterative thresholding [27] can work well if the cells are isolated and foreground and background intensity ranges are sufficiently far apart. Combined with morphological filtering, they enable a fast binarization [29].

Graph cut based binarization in the context of cell nuclei has been applied several times. Daněk [8] heuristically determines hard constraints for the data term, the boundary term employs a sophisticated computation of the distance between pixels in a gradient induced Riemannian manifold. This technique has high computational demands. Al-Kofahi et al. [1] perform minimum error thresholding to gain probability distributions for the data term, and use a simple gradient-based smoothness term.

If the cells are isolated, a single image binarization can be used to successfully segment the nuclei. However, for touching or overlapping objects, the task is more complex. The binarization step computes an *under-segmentation* of the image, resulting in clusters of cell nuclei. These have to be split in a cluster separation step.

1.1.3 Seed Detection

In terms of image analysis, detection refers to the task of placing a marker, or a seed point, on each desired object in the image. A basic insight of image analysis is that detection is easier than segmenting, i.e. delineating the border of objects. Therefore, an actual segmentation is often preceded by a seed detection phase. This especially applies to the field of cell nuclei segmentation, where the desired objects are roundish *blobs* of similar sizes.

The popular multiscale Laplacian of Gaussian (LoG) provides an effective blob detector. Stegmaier et al. [45] find peaks in the maximum intensity projection of a range of predefined scales, while Al-Kofahi et al. [1] implement a computationally more expensive automatic scale detection, exploiting cues in the distance transform of the binary foreground mask. Daněk and Matula [9] search for peaks in the distance transform. While this method is fast, it is error prone for tightly clustered nuclei, especially in 3D. Lou et al. [33] use a blob detector based on eigenvalues of the Hessian matrix.

The success rate of later stages depend on the accuracy of the seeds, as each seed will later form a final object.

1.1.4 Cluster Separation

When faced with an under-segmentation, object clusters have to be split to distinguish each cell nuclei. The classical approach is the watershed algorithm. However, it often results in over-segmentation, making a post-processing step necessary [29]. The graph cut framework with its multi-label variant has been used as well. Daněk and Matula [9] use the above described seed points as initial labels and combine the boundary term from the binarization phase with a distance transform induced shape term. In [1], the data term is modeled with a Gaussian mixture for each cell, and the ad hoc smoothness term is reused. The α-expansion technique is sped up by reducing the number of labels. Instead of introducing one label per seed, the labels are found by coloring a graph which is built on the initial labels. In [33] the probability distribution is gained with supervised learning, by training a random forest on a set of local features. This approach requires a set of images with ground truth – a fact which makes this method difficult to employ for huge 3D images with thousands of cell nuclei. The authors also introduce shape priors: they augment the graph cut formulation with a term which encourages cuts to be aligned to a vector field centered at the initial seeds.

He et al. [19] employ the normalized cut algorithm [42] to segment 2D cell nuclei images. The image is represented with a fully connected graph. The edge weights encode gradient information, and weights of edges which cross binarization contours are set to zero. The optimal normalized cut is computed by solving an eigenvector system. Clusters consisting of more than two nuclei are split recursively, with a heuristically fixed number of recursions. Depending on this number, the image might be over- or under-segmented. The complexity of solving the eigenvector system is $\mathcal{O}(n^3)$, where n is the number of image pixels. It is unclear if this method can be extended to big 3D images.

1.1.5 Cluster Merging

Instead of performing a top-down approach, where an under-segmentation is split up in a post-processing step, it is also possible to segment the image in a bottom-up fashion. The idea is to use an algorithm which results in an over-segmentation and then merge the components to intact nuclei. Lin et al. [29] apply a 3D watershed algorithm on a gradient-weighted distance transform, which allows them to incorporate both geometric and intensity cues. In a postprocessing step, they recursively merge touching objects. Merging blocks of pixels is conceptually simpler than splitting them, because the cut is already defined. The difficulty is to determine which segments are to be merged. In [29], a score is computed which signifies the coherence of the potential merged objects to a nuclei model. The highest score determines which objects are merged next. However, future potential merging steps are not considered. An improvement Lin et al. [30] is to make merge decisions by finding the best solution in merge trees built on the nodes of a region adjacency graph. This graph encodes all possible sequences of merging steps. Liu and Sclaroff [32] combine splitting and merging of objects in one algorithm. The algorithm recursively splits regions by straight lines between candidate endpoints.

1.1.6 Prior Knowledge and Model-driven Approaches

Best segmentation results cannot be obtained with an all-purpose algorithm. The assumption that the information required for a perfect segmentation is provided by the image only is wrong. Hence, segmentation algorithms have to incorporate domain knowledge to be successful. This knowledge can be of qualitative or quantitative nature. In the context of nuclei segmentation, for example, the qualitative knowledge that cell nuclei are blob-shaped is exploited by using a LoG seed detector. Quantitative knowledge is introduced by making use of known values for certain object-specific features. In Stegmaier et al. [45], for example, the multiscale LoG-filter works on a range of user-defined scales, which reflects the expected diameter range of the nuclei.

The knowledge on a feature or a set of features can be expressed in a probabilistic *object model*. This model can guide a segmentation process by reflecting the probability that an image object is a desired object or not. Instead of predefining the model parameters, it is also possible to learn them from the data. The merge decision process in [29, 30], for example, makes use of a model built on eight features, including texture, volume and several quantifiable shape features. The parameters of a multi-dimensional Gaussian distribution over these features are estimated from heuristically chosen objects in the image. This distribution enables the computation of a confidence score for any object. Liu and Sclaroff [32]

The split and merge algorithm in [32] is model-driven as well, the parameters for the three features color, shape and area overlap are gained by supervised learning.

1.2 Preliminiaries

1.2.1 Images

This section introduces a set of terms and their relation to each other. For a deeper explanation of these terms, we refer to [43, 16, 48]. A *scene* is a segment of the real world. It is composed of three-dimensional *objects*. A *real image* is a continuous function $f : \mathbb{R}^2 \to \mathbb{R}$. It maps points from the two-dimensional *image plane* to the *intensity*. The 2D real image is the result of a projection of the scene onto the plane [43]. To be processed by a computer, the real image is *digitized*. A *digital image* I is a function from a bounded discrete rectangle $\mathcal{P} =$

$\{0, 1, \ldots, S_x\} \times \{0, 1, \ldots, S_y\}$ to the discrete brightness levels $V = \{0, 1, \ldots, K\}$. An element $p = (i, j) \in \mathcal{P}$ is called a *pixel*. The rectangle $\mathcal{P}$ is called the *pixel space*. For the value of a pixel p, we also write I_p instead of $I(p)$.

The digitization is the conversion of the real image to the digital image. The discretization of the plane is done by *sampling* with a discrete set of sampling points in the plane. *Quantization* splits the intensity range into K intervals. Usually, K is of the form 2^b, where b is the *bit depth*. The *histogram* $h : V \to \mathbb{N}$ of a digital image is the absolute frequency of discrete brightness levels in an image.

Digital images formed by the described process are two-dimensional, because the image sensor used in the digitization is two-dimensional. Images of higher dimensions are formed by grouping a set of images to an *image stack*: a three dimensional image is a stack of two-dimensional images. Multiple 3D images which represent different points in time can be stacked to a 4D (or 3D+t) image. This work, only deals with 3D images. In this case, the pixel space is a discrete cuboid:

$$\mathcal{P} = \{0, 1, \ldots, S_x - 1\} \times \{0, 1, \ldots, S_y - 1\} \times \{0, 1, \ldots, S_z - 1\}.$$

A 3D pixel might also be called *voxel*. The *size* S of an three-dimensional image is the vector (S_x, S_y, S_z). In the following, when we speak about images, we will refer to 3D digital images.

There are two notions of distances in digital images: *physical distance* or *pixel spacing*, and *pixel distance*. To be able to measure the physical distance between two real image features using the digital image, the pixel spacing has the be known. It is defined as the shortest distance between two real-world points which can be distinguished with the combination of the optical system and the employed camera, along each of the three axes. Given the image spacing $\delta \in \mathbb{R}^3$, $\delta = (\delta_x, \delta_y, \delta_z)$, the image *resolution* denotes the pixel per unit distance $r = (r_x, r_y, r_z) = (1/\delta_x, 1/\delta_y, 1/\delta_z)$. Then the physical distance $d(p, q)$ between two pixels $p, q \in \mathcal{P}$ is computed by interpreting $\mathcal{P}$ as the euclidean three-dimensional space and determining the euclidean distance between p and q.

The pixel distance is a dimensionless measure for the amount of basic steps required to move between discrete pixels. We define two variants. The distance D_6 is the 3D extension of the city-block distance and is defined as

$$D_6((x_0, y_0, z_0), (x_1, y_1, z_1)) = |x_1 - x_0| + |y_1 - y_0| + |z_1 - z_0|.$$

D_{26} is known as the checker-board distance:

$$D_{26}((x_0, y_0, z_0), (x_1, y_1, z_1)) = \max(|x_1 - x_0|, |y_1 - y_0|, |z_1 - z_0|).$$

Two pixels p, p are *k-adjacent*, $k \in \{6, 26\}$, if $D_k(p, q) = 1$. Then, p is a k-neighbor of q. The *k-neighborhood* of a pixel is the set of pixels which are k-adjacent to it. When clear from the context, or not necessary, we omit the specifier in the following. A *neighborhood system* $\mathcal{N}$ is the set of adjacent pairs of pixels.

A *path* from a pixel p to a pixel q is a sequence of pixels $A_1, A_2, \ldots, A_n$ where $p = A_1$, $q = A_n$ and A_{i+1} is a neighbor of A_i for $1 \leq i < n$. Two pixels are *connected* if there exists a path between them. A *connected component* or simply *component* is a set of pixel which are pairwise connected. The *border* of a component C is the set of pixels of the component with one or more neighbors outside of C.

A *segmentation family* of an image is[2] a finite set of subregions $R_1, \ldots, R_l$, where

- R_i is a connected component,
- $\mathcal{P} = \bigcup_{i=1}^{l} R_i$, and
- $R_i \cap R_j = \emptyset$ for $i \neq j$.

It is possible to relax the second demand by introducing a special not necessarily connected subregion R_b for the background defined as $R_b := R \setminus \bigcup_{i=1}^{l} R_i$.

For a given segmentation family, a *segmentation* is a function which maps pixels to the subregions they are contained in. The subregions R_i are then also called *labels*. A special case of a segmentation is a *binarization*, where the number of subregions is $l = 2$, and the regions are foreground and background, denoted as R_f and R_b, respectively.

Informally, each component of a segmentation corresponds to a real world object in the scene represented by the digital image. Hence, each pixel is mapped to either a real world object, or background. Another notion of segmentation is *soft segmentation*, which takes into account that a particular pixel can incorporate signals of multiple objects. This is accomplished by making use of fuzzy set theory [50]. Each pixel is mapped to a fuzzy set over $R_1, \ldots, R_s$. This technique should not be confused with the uncertainty propagation discussed in the next section, where fuzzy sets are used to measure the uncertainty of the entire (hard) segmentation. We only consider hard segmentation in our work.

1.2.2 Uncertainty Evaluation

Most operators in an image processing pipeline are unable to give precise results for all inputs. On the one hand, this is due to the inherent deficiency of the imaging quality. The low signal-to-noise ratio might it make hard to distinguish small nuclei from background, or aberrations might be wrongly detected as nuclei. On the other hand, many tasks in image processing are ill-posed problems [28, 8]: the solution to such a problem are subject to ambiguities. There might not be exactly one correct answer regarding a certain image feature, e.g. a cell which is in the mitosis during the imaging process can be detected as a single cell or two cells, depending on how far the mitosis is already visible. In order to prevent errors made by one operator from affecting downstream tasks, but also to avert the loss of improbable but viable solutions, a stage should be able to inform its successors of the uncertainty involved in its output. This is made possible by introducing a measure for the validity for every extracted piece of information and propagating it together with the results to the next pipeline stage [44].

We follow Stegmaier et al. [44] in making use of fuzzy set theory to measure uncertainty, based on *a priori* knowledge. They introduce a fuzzy set membership function $\mu_{ijm} \colon \mathbb{R} \to [0, 1]$ for every operator i, feature m and linguistic term j. One way to define such a function is by using the following set of four terms: "Feature is... {... as expected";... too small, but useful";... too large, but useful";... not useful"}. This set lets us model the uncertainty with a trapezoidal membership function $\mu_m(x, \theta_m)$ defined by a quadruple $\theta_m = (a, b, c, d)$ as

[2] Most definitions[16] further demand that $Q(R_i) =$ `TRUE` and $Q(R_i \cup R_j) =$ `FALSE` for a logical, domain-specific predicate Q. While this is certainly true for all segmentations, it is rather vague not necessary for a definition.

$$\mu_m(x, \theta_m) = \begin{cases} \frac{x-a}{b-a} & \text{if } a \leq x < b; \\ 1 & \text{if } b \leq x < c; \\ \frac{d-x}{d-c} & \text{if } c \leq x < d; \\ 0 & \text{otherwise.} \end{cases}$$

Following the rules of fuzzy set theory, multiple uncorrelated features can be combined by multiplying their respective membership functions.

1.2.3 Balanced Graph Partitioning

Given an undirected graph $G = (V, E)$ with non-negative edge weights $c : E \longrightarrow \mathbb{R}^+$, and a number $k \in \mathbb{N}$, $k > 1$, the *graph partitioning problem* is to find k sets of nodes $V_1, \ldots, V_k$, with $\bigcup_{i=1}^{k} V_i = V$ and $V_i \cap V_j = \emptyset \ \forall i \neq j$. Given an additional imbalance factor $\varepsilon \in \mathbb{R}^+$, the *balanced graph partitioning problem* further demands that the sets are balanced, i.e. that $|V_i| \leq (1 + \varepsilon)\lceil |V|/k \rceil$. The goal is to minimize an objective function. In this work, we will aim to minimize the sum of the weights of the set of cut edges $\{\{u, v\} \in E \mid u \in V_i, v \in V_j, i \neq j\}$. The graph partitioning problem has been shown to be NP-complete [22, 14]. Practical graph partitioning tools rely on heuristics to find good partitions. Most tools use the multilevel approach [20], where the graph is first iteratively coarsened, i.e. a hierarchy of graphs is constructed which keeps the original structure but reduces the input size. A common method for coarsening is contracting edges in a matching in G [20]. Then a partitioning is computed on the coarsened graph. One way to do this is to search for k seed nodes which are far away from each other and then alternately run k BFS, one from each seed [12]. Finally the graph with its partitioning is unpacked, while employing local search methods, such as max-flow min-cut based search [40] and FM [13]. As we cannot cover all methods here, we refer the reader to [2] for a deeper insight. From the many available tools and libraries for graph partitioning, we tested METIS [26] and KaHIP [41], and achieved best results with a fine-tuned configuration for KaHIP.

compute exact solutions to the two-label segmentation problem, i.e. image binarization.

1.2.4 Graph Cut Segmentation

The graph cut framework is a powerful graph-based segmentation method. It has been shown [17, 6, 3] that it is possible to efficiently minimize an *energy function* of the form $E(A)$ for a binarization A defined as

$$E(A) = \lambda E_{\text{data}}(A) + E_{\text{smooth}}(A).$$

A minimum of this function corresponds to an exact solution to the two-label segmentation problem, i.e. image binarization.

The data term is typically

$$E_{\text{data}}(A) = \sum_{p \in \mathcal{P}} D_p(A_p),$$

where D_p measures the pixel-wise penalty of mapping pixel p to label A_p, i.e. fore- or background. The smoothness term $E_{\text{smooth}}(A)$ penalizes neighboring pixels with different labels. Using a neighborhood system $\mathcal{N}$, it is defined as

$$E_{\text{smooth}}(A) = \sum_{(p,q) \in \mathcal{N}} B_{p,q} \cdot 1_{A_p \neq A_q}.$$

Here, $1_{A_p \neq A_q}$ is an indicator function which is 1 iff $A_p \neq A_q$. The function $B_{p,q}$ is the boundary energy. The energy function $E(A)$ is minimized by computing a minimum cut on a special graph.

Boykov et al. [6] extended the method to multi-label energy minimization, which is NP-hard in general. They find approximate solutions by introducing two approximation techniques, α-expansion, and α/β-swap. The former performs iterative binary segmentations between one label and every other label. The latter computes the minimum cut between all pairs of labels.

The graph cut segmentation method is only a framework. To employ it in practice, the terms D_p and $B_{p,q}$, and the scalar λ have to be parameterized. Boykov and Funka-Lea [3] suggest to determine models for the pixel intensity distributions of the objects to be segmented. The data term can then be parameterized using the fit of the pixel value to the models. Hard constraints can be incorporated as well.

Due to the extensive use of the graph cut framework for segmentation problems, maximum flow algorithms which exploit the special structure of the graph have been developed [5, 15] and parallelized [10, 31, 25].

1.2.5 Cut Metrics

Due to the discrete nature of digital images, computing the surface area of a real-world object using a digital 3D image[3] is a nontrivial problem. Interpreting each voxel as a cuboid and summing up the area of their surfaces is only a bad approximation, since with that method, every object has the same surface area as its bounding box. In the context of adopting graph cut to compute geodesics in Riemannian spaces, Boykov and Kolmogorov [4] have introduced the notion of *cut metrics*. The authors have shown that it is possible to compute edge weights for a regular grid graph such that the length of a contour is approximated by the sum of the edge weights cut by that contour. Their method is based on the Cauchy-Crofton formula. This formula relates the Euclidean length $|\mathcal{C}|$ of a curve $\mathcal{C}$ to the function $n_{\mathcal{C}}(l)$ which measures the number of intersections between the curve $\mathcal{C}$ and a line $l \in \mathcal{L}$:

$$|\mathcal{C}| = \frac{1}{2} \int_{\mathcal{L}} n_{\mathcal{C}}(l) \, \mathrm{d}l = \frac{1}{2} \int_0^{\pi} \int_{-\infty}^{+\infty} n_{\mathcal{C}}(\phi, \rho) \, \mathrm{d}\phi \, \mathrm{d}\rho, \tag{1}$$

using that the space of lines can be represented as $[0, \pi] \times [-\infty, +\infty]$. Boykov and Kolmogorov [4] approximate the integral with partial sums by partitioning the space of lines among the edges of the grid graph's neighborhood system. This method yields edge weights

$$\omega_k = \frac{\delta\phi_k \delta\rho_k}{2}, \tag{2}$$

where $\delta\phi_k$ is the angular partition of edge k and $\delta\rho_k$ is the distance of the line induced by edge k of the neighborhood system to the nearest parallel edge induced line.

The method can be formulated for 3D grids as well. However, for that case Boykov and Kolmogorov [4] do not specify how to partition the space of angular orientations among the edges of the neighborhood system. Daněk and Matula [9] proposed a solution both in 2D and 3D, as well as for grids with anisotropic node spacing. They compute the Voronoi diagram of the intersections of the neighborhood system with a unit hypersphere. The angular partitioning of an edge is equal to the fraction of its Voronoi cell.

[3] or the length of a contour in 2D

2 Our Method

This section deals with our segmentation method. Our method performs a top-down approach. We first perform a fast and robust foreground detection, detailed in Section 2.1. The connected components in the binarized image are then recursively split. We make use of balanced graph partitioning to cope with the bias towards small partitions when minimizing a cut (cf. Section 2.2). Our segmentation framework computes a score for each component which is based on prior knowledge. The nucleus model is presented in Section 2.3 and the segmentation framework Section 2.4. The advantage of our cluster segmentation approach over the approaches described in Section 1.1 is that our method does not depend on a heuristic seed detection, which improves the detection and segmentation quality for highly clustered cell nuclei.

2.1 Binarization

The first step of our segmentation algorithm is to binarize the image. We aim to classify the image pixels into two groups, the foreground and the background. This step immediately reduces the input, as all background pixels do not have to be further examined. For our purpose, it is important that foreground objects not be wrongly classified as background. An error of this kind leads to undetected nuclei because our subsequent steps are not capable of detecting these objects. Another binarization mistake are non-foreground areas which are wrongly detected as foreground. This case might arise due to defective camera sensors, e.g. isolated white pixels. To some extend, our splitting framework is able to detect and discard these erroneous regions especially in the aforementioned case.

A third binarization error is an inexact delineation of the foreground object's border. The object is correctly detected as foreground, but its shape or size differs. This error commonly arises in binarized images of confocal fluorescence microscopy. On the one hand, unstained parts of a nucleus may appear as dark as the background [47]. On the other hand, an enlargement of the nuclei can be observed for binarized cell nuclei images. The diameter of binarized nuclei is several pixels larger than the real extent of the nuclei. The reason for the inflated objects is the optical diffraction, which can be expressed by the *point spread function* (PSF) of the imaging system. The PSF represents the response of a pixel sensor to an ideal point light source [48]. A point source appears blurred in the image and might be surrounded by concentric rings. The area directly adjacent to the foreground objects is then still brighter than the dark background. This effect can add up if several objects are close to each other. In these cases, clustering-based thresholding approaches which do not use domain knowledge result in inflated objects.

This phenomenon effects our nuclei segmentation method both in quality and in running time. Firstly, the quality suffers. Our method relies on splitting the foreground regions of the image. It only adds cuts inside the foreground, but does not change the delineation between fore- and background. Furthermore, the inflation of the nuclei can also cause well-separated objects to be connected in the binarization. These connecting bridges will then be part of one the segmented objects. These circumstances would have a negative impact on the quality of the segmentation results. Secondly, the segmentation speed would be slowed down. Our method uses graph partitioning to split the foreground objects. The larger number of merged objects requires more partitioning steps, and their inflated size results in overall bigger graphs to be partitioned.

The PSF depends on the imaging system. It is a combination of the blur induced by the optical system and the finite integration area of a chip sensor [48]. Hence, the extend of the

foreground blur can differ across distinct datasets. In one dataset, the blur might be strong, and measures have to be taken to detect and cope with it. In another set, simple thresholding already leads to good results. In this work, we experimented with three different binarization methods. The first is Otsu's method [36], a simple thresholding algorithm. The other two are based on work by Restif [37], who models the gray level histogram of a blurred image. The model is described in the next section. We either use this model to compute a threshold, or parameterize the graph cut framework with its probability distributions.

For all algorithms, binarization quality could be improved by convolving the image with a low-pass Gaussian filter with small standard deviation σ_s. This method reduces background noise and eliminates one-pixel errors. Another image degradation specific to confocal microscopy is uneven illumination, a brightness gradient in the image, often along the axial direction [35]. We overcome this deficiency by dividing the image into sets of consecutive slices of size m and binarizing each set separately.

2.1.1 Image Histogram Model

We compute the normalized gray level histogram $h(i)$ with the goal to infer two distributions, one for the foreground and one for the background. Al-Kofahi et al. [1] assumed both to be Poisson distributed, citing the image formation process and their findings of the histograms to be bimodal. However, for some of our data we observed unimodal histograms, due to the small fraction of foreground compared to the entire image, the high intensity variance in the fluorescing nuclei, and the system-specific point spread function, which blurs the image and results in actual background regions to be brighter around the foreground objects. We therefore adopt a model inspired by [38]. The background is subdivided into non-illuminated (NB) and illuminated background (IB). The former is a normal distribution with parameters μ_b and σ_b, and a priori probability p_b:

$$NB(i) = \frac{p_b}{\sqrt{2\pi}\sigma_b} \exp\left(-\frac{(i-\mu_b)^2}{2\sigma_b^2}\right). \tag{3}$$

The IB accounts for the fluorescence microscopy-inherent blurring effect: an increased gray level intensity of the background around the fluorescing nuclei. This blur is modeled with an exponential decline, from which the model for the IB is derived (see [38]) as

$$IB(i) = \frac{2\alpha A}{i-\mu_b} \log \frac{(I_f - \mu_b)}{(i-\mu_b)}, \tag{4}$$

where I_f is the intensity at the border of the object and A is the object's area. The factor α defines the blur's rate of decline. Restif [38] introduces a foreground normal distribution and an IB term for each nuclei, which is practical as the images used in his work on average only contain about seven foreground objects. For our purpose, we model the nuclei with one normal distribution $F(i)$ with parameters μ_f, σ_f, and p_f, and the entire illuminated background with one term $IB(i)$. We set I_f to μ_f, and A to p_f, the foreground's a priori probability. Because the IB encroaches into the NIB, we only compute the $IB(i)$ in the range $[\mu_b + 2\sigma_b, \mu_f)$. Then, the histogram model is

$$h_{model}(i) = NB(i) + IB(i) + F(i) \tag{5}$$

$$= \frac{p_b}{\sqrt{2\pi}\sigma_b} \exp - \frac{(i-\mu_b)^2}{2\sigma_b^2} + \frac{2\alpha p_f}{i-\mu_b} \log \frac{(\mu_f - \mu_b)}{(i-\mu_b)} + \frac{p_f}{\sqrt{2\pi}\sigma_f} \exp - \frac{(i-\mu_f)^2}{2\sigma_f^2} \tag{6}$$

with the set of missing parameters $\{p_b, \mu_b, \sigma_b, p_f, \mu_f, \sigma_f, \alpha\}$.

Like Restif [38], we determine the parameters with the expectation–maximization (EM) algorithm [11]. The EM algorithm iterates between an expectation step (E-step) and a maximization step (M-step). The former computes weights which define the proportion of each summand in Equation 5, given an estimate of the parameters. The latter estimates the parameters, given a set of weights.

We introduce weights $\omega_i^{nb}, \omega_i^{ib}, \omega_i^f$ for the non-illuminated background, illuminated background, and foreground, respectively. In the E-step, we use Equation 5 to compute the weights for each histogram bin i, e.g.

$$\omega_i^b = NB(i)/h_{model}(i) \tag{7}$$

In the M-step, the parameters for the normal distributions are estimated based on the weighted histogram:

$$p_b = \sum_i \omega_i^b h(i), \tag{8}$$

$$\mu_b = \sum_i i\omega_i^b h(i), \quad \text{and} \tag{9}$$

$$\sigma_b = \sqrt{\sum_i (i-\mu_b)^2 \omega_i^b h(i)}, \tag{10}$$

for the background distribution and equivalently for the foreground distribution. The parameter α is subsequently determined, analogous to [38], as

$$\alpha = \frac{1}{\log \varepsilon} \sqrt{\frac{1}{p_f} \sum_{i=\mu_b+2\sigma_b}^{\mu_f - 1} \omega_i^{ib} h(i)}, \tag{11}$$

where $\varepsilon = 2\sigma_b/(\mu_f - \mu_b)$ is a correction factor for the truncation of the illuminated background model.

Alternating E-step and M-steps, the algorithm converges after a small number of iterations (about 5 in our settings). We found the outcome of the EM algorithm to be robust and largely independent on the initialization. Therefore, we initialize the parameters with a simple iterative thresholding algorithm [39], where a threshold is defined as the histogram mean, and a new threshold is iteratively computed as the average between the means of the histogram below and above the old threshold. We use the final threshold to compute the means and a priori probabilities of the foreground and background distribution, and set the variances to 1 and α to 0.01.

For subsequent computations which require probability distributions for foreground and background pixels, we will combine the illuminated and non-illuminated background into one term $B(i)$:

$$B(i) = NB(i) + B(i). \tag{12}$$

Hence, $B(i)$ ($F(i)$) can be interpreted as the probability that a pixel in a background (foreground) area has gray level i.

We implemented two binarization methods using the above distributions. The first is to compute an intensity threshold $t^\star$ as

$$t^\star = \min_{i \in V} F(i) < B(i). \tag{13}$$

All pixels with values above this threshold are determined as foreground, all others are background.

The second method uses these distributions to parameterize the graph cut framework. With the notations in Section 1.2.4, we set the pixel-wise data term to

$$D_p(A_p) = \begin{cases} -\log F(I_p) & \text{if } A_p = R_f; \\ -\log B(I_p) & \text{if } A_p = R_b. \end{cases} \tag{14}$$

For the smoothness term, we use

$$B_{p,q} := \exp\left(-\frac{(I_p - I_q)^2}{2\sigma_{\text{bin}}^2}\right) \frac{1}{d(p,q)}. \tag{15}$$

This function was suggested in [3] and also adopted in [1]. However, it introduces a parameter, σ_{bin}. It results in a high penalty for discontinuities when the intensity difference is below σ_{bin}. A the function resembles a Gaussian probability distribution, the parameter can be interpreted as the standard deviation of the image noise. This parameter, and the scalar parameter λ, have to be set by hand.

2.2 Cluster Separation

The binarized foreground in cell nuclei images decomposes into a large number of small connected components. These can either form single nuclei, clusters of nuclei or falsely detected regions due to imaging defects. We perform graph-based segmentation successively on each component.

We use balanced graph partitioning to split the binarized foreground. The graph partitioning aims to minimize the cut between partitions, i.e. the sum of weights of the cut edges. This allows the method to take advantage of intensity or gradient information of the original image, but also of shape cues of the binary foreground mask.

In this perspective, the approach is related to the multi-label cut methods from the graph cut framework, namely α-expansion, and α/β-swap. However, these methods rely on a proper parameterization to show good performance. The data term of the labeling energy graph cut framework reflects the attractions of a voxel to all labels. These have to be determined in advance. With graph partitioning, the number and, more importantly, the positions of the labels do not have to be known. Another difference is that the graph cut methods perform multiple iterations of optimization steps. The α-expansion, for example, iterates over all partitions and computes a minimum cut to the rest of the graph. It repeats this process multiple times until convergence. In our method, each graph partitioner call is only done once, and its outcome is either kept or discarded.

The balance constraint serves multiple purposes. Most importantly, it counteracts the bias towards unnaturally small components. Without this measure, the process would almost always cut only a small set of pixels from the main component, as these are connected by a small number of edges. Figure 1a gives an example for this problem.

The balance constraint forces the graph partitioner to select edges in the green area of Figure 1b and thus finds the cut which correctly segments the image. This image also illustrates another reason to employ balanced graph partitioning for cell nuclei segmentation. The objects are often roughly equally sized and separated only by tight bottlenecks. Therefore, the combination of minimizing a cut and balancing the two resulting partitions makes adequate use of

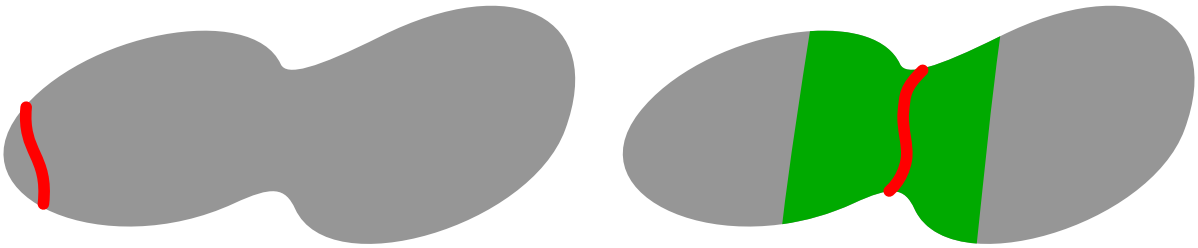

(a) Minimum cut bipartition of the component.

(b) Minimum cut balanced bipartition of the component. To obey the maximum imbalance factor, cuts have to lie in the green area.

Figure 1: Purpose of the balance factor for graph partitioning.

the shape cues of the binary foreground mask: cuts are smaller when they contain less edges, which favors bottleneck regions to be cut.

The balancing also improves the speed of our algorithm for big components which have to be split multiple times. While the number of required partitioner calls remains asymptotically unchanged, the subcomponents sizes decrease faster, and the recursion depth is decreased.

For each connected component in the binarized image we construct a graph with nodes representing voxels and edges connecting neighboring voxels. A 6-neighborhood was sufficient in our experiments. As the number k the foreground components have to be split into is not known a priori, we set $k = 2$ and compute recursive bipartitions. The model-based stopping-criterion will be explained in the next section.

2.2.1 Edge Weights

We describe the edge weights for the isotropic case. We implemented two methods of setting the edge weights. One (*grad*) sets weights based on gradient, the other (*prob*) on the probability of a pixel intensity belonging to background.

For the gradient-based edge weights, we reuse the method for the graph cut binarization (see Section 2.1). The intuition behind this choice is that neighboring pixels with high intensity difference are less likely to belong to same label. We compute edge weights $c_{\text{grad}}(\{p, q\})$ between two pixels p, q with intensities I_p and I_q as

$$c_{\texttt{grad}}(\{p, q\}) = \exp\left(-\frac{(I_u - I_v)^2}{2\sigma_{\text{grad}}^2}\right).$$

The parameter σ_{grad} is chosen smaller than the σ_{bin} in the binarization.

However, in our case all pixels are already detected as foreground. We experienced problems with this method for images with high variations of the fluorescent marker. The resulting irregular inner-cell texture of the nuclei can lead to wrong cuts if the inner-nuclei gradient is too high.

In contrast to 2D images, tightly clustered cell nuclei in 3D images do not overlap. In our instances, we observed that even between very close or touching nuclei there was still a tight band of darker pixels. Therefore, a second method is to induce the edge weights directly from the gray levels. The cut should tend to be near the darker area between the bright nuclei. In

order to be able to quantify the notion of "dark areas", we reuse the histogram image model described in Section 2.3.

We maximize the probability that the nodes a found cut S is adjacent to represent background pixels. To this end, we disregard the non-independence of neighboring pixel values, and aim to maximize the product of the background probabilities of the adjacent pixels. Formally, we maximize

$$\prod_{\{u,v\}\in S} \min\{\mathrm{P}(B|I(u)), \mathrm{P}(B|I(v))\},$$

where we use the conditional probability $\mathrm{P}(B|I(u))$ to denote the probability that pixel u with intensity $I(u)$ is in the background. Maximizing the above expression is equivalent to minimizing

$$\sum_{\{u,v\}\in S} -\log \min\{\mathrm{P}(B|I(u)), \mathrm{P}(B|I(v))\}.$$

From the last formula, we derive the edge weights $c_{\mathrm{prob}}(\{u,v\})$ as

$$c_{\mathrm{prob}}(\{u,v\}) = -\log \min\{\mathrm{P}(B|I(u)), \mathrm{P}(B|I(v))\}. \tag{16}$$

The required probabilities that a given intensity represents background can be computed with

$$\mathrm{P}(B|I(u)) = \frac{p_b \mathrm{P}(I(u)|B))}{p_b \mathrm{P}(I(u)|B) + p_f \mathrm{P}(I(u)|F)}. \tag{17}$$

for our model described in the previous section, this is equal to

$$\mathrm{P}(B|I(u)) = \frac{B(i)}{h_{model}(i)}. \tag{18}$$

For experimental purposes, we also tried setting all edge weights to the same constant. Then the cut is independent of the intensity levels and only depends on the shape of the foreground objects.

Once the graph is constructed, we partition it into two blocks with a balanced graph partitioning solver. In general, these blocks do not have to be connected. This circumstance occurs in particular if a component consists of more than two nuclei of differing sizes, as depicted in Figure 2a. Neither the red nor the green cut are eligible. One of the smaller cell nuclei would form one block, and the other block would consist of the big nuclei paired with the other small nuclei, which is not feasible if the imbalance factor is too tight. Instead, the balanced graph partitioning results in the green cut shown in Figure 2b. It partitions the graph into two blocks, one (lilac) of which is not connected. To cope with this phenomenon, we perform a breadth-first search after the partitioner call to identify the connected components and their sizes. In the case shown in the example, the component is partitioned into three subcomponents.

In order for the cut in Figure 2b still being valid, even for equally sized nuclei, the balance constraint has to be relaxed. Hence, we set $\varepsilon = 0.5$ to cope with this case.

2.2.2 Anisotropic Resolution

In most confocal fluorescence microscopy images, the resolution along the lateral axes (r_x and r_y) is smaller that the axial resolution (r_z). When setting the edge weights, we take the anisotropic resolution into account by dividing the weight of an edge between two pixels u and v by their distance $d(u,v)$.

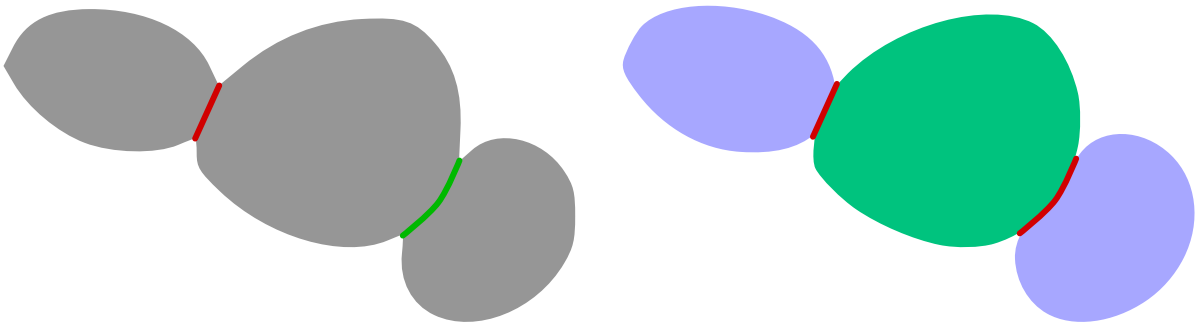

(a) Neither the red nor the green cut obey the imbalance factor for a bipartitioning.

(b) The minimum cut balanced bipartitioning. One of the two blocks (lilac) is not connected. A BFS results in three subcomponents.

Figure 2: Unconnected blocks of the graph partitioning.

This method is based on the observation that the cut of an object should be invariant of its orientation in the image. The number of edges cut by a 2D plane in the 3D image depends on its orientation. For an example, a cut of the form of an xy-plane of area 1 will cut (i.e. is perpendicular to) $r_x r_y$ edges, while a xz-plane of the same size will cut $r_x r_z$ edges. By dividing the edges by their length the sum of these cut edges will be the equal.

2.3 Nucleus Model

We construct a mathematical model of a cell nucleus. The goal is twofold: on the one hand, the model should decide whether an image component represents a single object such that the recursion can be stopped, a cluster of multiple objects to be split, or background noise to be discarded. On the other hand, we want to compute a score $s(C_i)$ for each component C_i which signifies the probability that it represents a single nucleus.

A splitting step results in a number of subcomponents, each of which is smaller in size than the original component. If a connected component is too big to be a nuclei, we can safely perform a splitting operation. If it is too small, we can be certain that it is not a nuclei. Therefore, in a first approach, we base the scoring function only on the volume V of the components. The volume V_i of an object C_i is approximated by multiplying its voxel number n with the spacing:

$$V_i = n\delta_x\delta_y\delta_z. \tag{19}$$

The voxel number n is counted during the breadth-first search.

From user-defined minimal and maximal cell volumes ($V_{\min}$ and $V_{\max}$, respectively) a trapezoidal fuzzy set membership function μ_V indicates whether a component has a plausible volume. We define μ_V with the parameter vector $\theta_V = (V_{\min}, (1+\lambda)V_{\min}, (1-\lambda)V_{\max}, V_{\max})$. The parameter λ gives us a fine control over the shape of the lateral sides of the trapezoid. In our experiments, it is set to 0.2.

A component with volume below V_{min} is discarded. A component is repartitioned if it is big enough to contain a minimum sized nucleus. Due to the imbalance factor ε of balanced graph partitioning, the reduction in size of a subcomponent compared to its original component is at least $(1+\varepsilon)/2$. Therefore, if V_i is greater than $\frac{1+\varepsilon}{2}V_{\min}$, the component C_i could be constituted of one minimum sized nucleus C_c. The remaining part, $C_i \setminus C_c$, does not have to represent a nucleus.

The volume can be used as a single criterion if the objects' sizes have a small variance. However, we observed that the volume range of cell nuclei allowed both a parent component and its subcomponents to be inside the user-given size limits.

To improve the decision making process and give better estimates for the uncertainty values, we use the knowledge that our desired objects are cell nuclei, which can be seen as roundish, blob-like objects. To distinguish between an ellipsoid-like nucleus and one cut into two nearly equally sized partitions (i.e. similar to an ellipsoidal dome), we compute the sphericity Ψ_c defined [51] as the ratio between the surface of a sphere with volume V_i and the surface A_i of the object C_i.

$$\Psi_i = \frac{\pi^{\frac{1}{3}}(6V_i)^{\frac{2}{3}}}{A_i}. \tag{20}$$

where A_i is the surface area of a C_i.

In theory, values for the sphericity Ψ are in a range between 0 and 1, where 1 is the sphericity of a perfect sphere. Depending on organism and cell type, different values of sphericity can be expected for cell nuclei in practice.

We assume an ideal nucleus to be an ellipsoid with near-equal semi-axes and deduced a sphericity of $\Psi_{ideal} = 0.96$ and $\Psi_{min} = 0.81$. The sphericity membership function is then defined as

$$\mu_\Psi(x) = \begin{cases} 0 & \text{if } x \leq \Psi_{min}; \\ \left(\frac{x-\Psi_{\min}}{\Psi_{\text{ideal}}-\Psi_{min}}\right)^2 & \text{if } \Psi_{min} < x < \Psi_{\text{ideal}}; \\ 1 & \text{if } \Psi_{\text{ideal}} \leq x; \end{cases} \tag{21}$$

We compute the surface area A with the cut metric results described in Section **??**. The set of cut metric edge weights in a regular grid are the same for each node. We use a 26-neighborhood for the cut metric evaluation because the approximation gained with a 6-neighborhood is to inaccurate.

We multiplicatively combine the two membership function values to gain the score $s(C_i)$ of a component:

$$s(C_i) := \mu_V(V_i)\mu_\Psi(\Psi_i).$$

2.4 Splitting Framework

We recursively split the connected components with balanced graph bipartitioning, computing the above score at each step. A way to see this splitting process is a depth-first search through a tree, a set of potential desired objects. Each inner node in the tree represents a component. The initial component forms the root of the tree, the children of a node are the outcomes of a partitioning step. The leaves are the desired objects. Listing 1 gives a top-level overview of this depth-first search framework.

The framework's core is a scoring function (ScoreFunction), which determinates the next processing step for a given component. A component can either be kept, discarded, or repartitioned. Keeping a component means that the search is stopped at that point and the component is not partitioned. In the DFS analogy the component is marked as a leaf. A component which is repartitioned might, but does not have to, consist of multiple desired objects. For a discarded component the search is stopped as well, but the component is deleted from the search tree. The scoring function also computes a score for each component.

```
1:  function RecursiveSplit(C) returns (k, R)        // k: number of components
2:      R ← PartitionerCall(C, k)                    // R: partitioning of C
3:      for each C_i ∈ R do
4:          (decision, s(C_i)) ← ScoreFunction(C_i)
5:          if decision = repartition then
6:              (k_i, R_i) ← RecursiveSplit(C_i)     // recursive call
7:              if k_i = 0 then                      // backtrack
8:                  if s(C_i) > 0 then decision ← keep
9:                  else decision ← discard
10:             else
11:                 merge R_i into C_i of R
12:         if decision = keep then k_i ← 1
13:         else if decision = discard then
14:             remove C_i from R
15:             k_i ← 0
16:     k ← ∑_i k_i
```

Algorithm 1: The cell partitioning framework.

The score $s(C_i)$ measured for a component C_i is interpreted as the probability for an event X_i, that the component C_i is a desired object, under the condition that its parent component C_p is not a desired object:

$$s(C_i) = \mathrm{P}(X_i \mid \bar{X}_p). \tag{22}$$

A component which is kept must have a score value greater than 0. If both the score $s(C_p)$ of a component C_p and the score $s(C_c)$ of one of its children is strictly positive, we keep the component which has a higher probability to be a nuclei. Using the fact that for the event X_c to occur, the event X_p cannot occur, we deduce:

$$\begin{aligned} & \mathrm{P}(X_c) > \mathrm{P}(X_p) \\ \Leftrightarrow \quad & \mathrm{P}(X_c \cap \bar{X}_p) > \mathrm{P}(X_p) \\ \Leftrightarrow \quad & \mathrm{P}(X_c \mid \bar{X}_p) \cdot \mathrm{P}(\bar{X}_p) > \mathrm{P}(X_p) \\ \Leftrightarrow \quad & s(C_c) \cdot (1 - s(C_p)) > s(C_p) \end{aligned} \tag{23}$$

Thus, for each child component of a component C_p, we check whether it is discarded or not with Equation 23. As an implication, components with a score greater than 0.5 are not split.

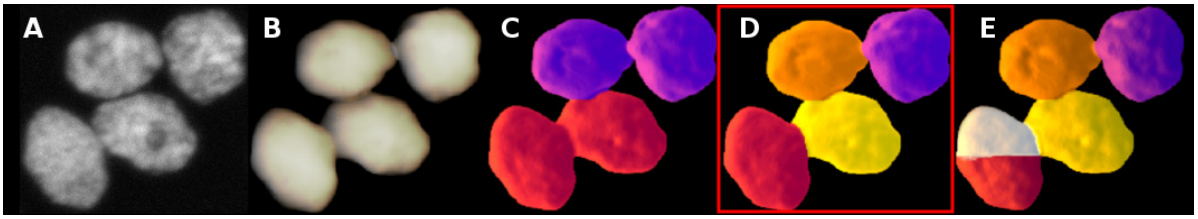

Figure 3: An exemplary execution of our pipeline. The input image (A) is first binarized (B), then recursively split (C-E). The split nuclei in the last image (E) is detected by our cell nucleus model. Image D is the final segmentation.

We examine each child separately such that discarded components have no influence on their siblings. That way, our algorithm is allowed to discard areas which do not contain any nuclei, reflecting the possible existence of wrongly binarized areas. In the backtracking step of our depth-first search, we check if any of the child components were kept. Otherwise, the current (parent) component is discarded or kept if its score is greater or equal to 0, respectively. Figure 3 gives an example. The purple nuclei in image D has a low score of 0.46 due to its upper-ranged size. It is split into two components (E), with scores 0.12 and 0.05. After comparing both to the original score, they are discarded during backtracking. Image D is the final segmentation.

3 Experiments

3.1 Setup

We implemented our algorithm in C++ using the Insight Toolkit [23] for image operations. The test machine is an Intel Core i7-920, equipped with four cores clocked at 2.67 GHz. The machine has 12 GiB of RAM. All experiments were performed on a single core. The graphs are partitioned with the open source graph partitioner KaHIP [41] with a customized configuration. We compare our algorithm to TWANG [45] and the graph cut (GC) segmentation described in [1]. We also include our binarization method (Otsu) in the evaluation.

We measured the quality of a segmentation by counting the number of added, split, merged and missing objects. Added and split are false positives, areas detected as nuclei where there is none. Merged and missing are false negatives. We divide by the number of labeled nuclei in the respective image. The percentages are then averaged over all images in the set. The times measured are averaged as well and do not include image I/O. None of the implementations were explicitly optimized for memory consumption. Nevertheless, we display values for memory consumption by extracting the peak resident set size of the respective process.

3.2 Instances

Ground truth of real world 3D microscopy images which contain thousands of cell nuclei does not exist. Manual segmentation of these images is not an option due to the large amount of objects on the one hand and the researcher's bias on the other hand. We therefore benchmark the quality of our algorithm on two sets of synthetic cell nuclei images. We downloaded a set of 30 HL60 cell line images generated with the CytoPacq toolbox [47], with high noise and clustering probability of 75% (High75). They contain 20 nuclei each and have an isotropic resolution. We scaled the images down by a factor of two in each dimension to let nuclei size resemble those of real world images. Each downscaled image consists of 64 slices with a resolution of 403×282. Based on this toolbox, [46] have conceived a more realistic set of 3D+t benchmark nuclei images. We use a set of 10 images of different time points (SBDE) which contain between 316 and 1016 nuclei. They have a size of $128 \times 640 \times 640$ and the physical distance between two slices is five times greater than the resolution of each slice.

3.3 Variants and Parameters

On the synthetic datasets, the best binarization results were achieved with applying Otsu's thresholding method after convolving the input image with a low-pass Gaussian filter with parameter σ_s. This procedure showed robust results on our image sets, even on those without

strictly bimodal intensity histograms. Computing a sophisticated image histogram model (Section 2.1.1) and using the parameters to threshold the image had worse results. Parameterizing the graph cut framework exhibited a bad trade-off between binarization quality and speed.

For High75, the partitioner edge weights using on the intensity probability (c_{prob}) had slightly better segmentation result than the gradient-based edge weights (c_{grad}). For SBDE however, the situation was inversed, with the edge weights computed by c_{prob} failing to capture the true object boundaries. The reasons are that the inner-nuclei variance is very low for SBDE, while the nuclei are very tightly clustered, such that there is no "background area" between the objects. We therefore use c_{grad} in our comparison.

Table 1: Parameters for High75 and SBDE, respectively

Parameter	High75	SBDE
σ_s	1.9	0.7
m	1	16
σ_{grad}	15	100
(V_{min}, V_{max})	$(20000, 39000)$	$(2900, 8550)$

Table 1 contains values for the parameters we used in our comparisons. The standard deviation σ_s is used for the Gaussian filter preceding the binarization and depends on the size of the objects to detect. The parameter m is the number of regions the image is split into for the binarization and reflects the change of the intensity distribution along the visual axis. The edge weights are parameterized with the parameter σ_{grad}, which separates the inner-nucleus intensity variance from the inter-nucleus intensity variance. The minimum and maximum volumes V_{min} and V_{max} are the limits used by our model. The difference in volume between the two image sets stems from the dimensionless nature of the synthetic images.

3.4 Comparison with other Methods

Table 2: Segmentation results for High75.

Algo	[s]	[MiB]	Added	Missed	Merged	Split
Otsu	0.4		0	0	71.3	0
GC	244	174	0	0	0.5	1.2
Ours	8	432	0	0	0	0
TWANG	14	251	0	0	0	0

Table 2 shows results for the High75 dataset. Both TWANG and our algorithm have no segmentation error on this dataset. GC has a low ratio of merged nuclei and split nuclei. Our algorithm is the fastest. With the unscaled images, quality is unchanged. GC takes three and a half hours per image, TWANG about 4 minutes and our algorithm less than 90 s, which shows that our algorithm scales best with increasing resolution. This result is due to the fast binarization, which can quickly (415 ms) reduce the amount of voxels to look at by about 91 %.

Table 3 lists results on the SBDE dataset. None of the tested algorithms exhibits perfect segmentation quality. The seed-based methods have very low rate of false positives, however

Table 3: Segmentation results for SBDE

Algo	[s]	[GiB]	Added	Missed	Merged	Split
Otsu	4		0	0	56.2	0
GC	147	1.02	0	6.0	3.3	0.2
Ours	56	0.83	0.4	0.8	1.3	1.8
TWANG	140	1.35	0.1	4.9	1.3	0

fail to detect at least 5% of the nuclei. The missing nuclei of GC are mainly at the far end of the visual axis, because GC has no mechanism to cope with the variations of illumination. The quality of our algorithm is similar or better than the other methods, only the number of split nuclei is elevated.

Our algorithm is again best regarding time. However, the running time varied between 10 s and 150 s. Running times of a segmentation algorithm can depend on several factors. The running time of our method is primarily affected by the number of clustered nuclei and the image resolution. The former is reflected in the number of partitioner calls, the latter increases the time needed per single partitioner call due to a larger graph. The running time variance of our method is caused by the variance in clustered objects for this data set, which ranges from 200 to 700 nuclei. Errors in an image analysis pipeline are inevitable. An approach to cope with this fact is to quantify the uncertainty in the outcome of an operator [44]. Our algorithm could be augmented by interpreting the score of our nuclei model as an uncertainty value. The error rate could then be reduced for narrow decisions of the nuclei model. A manual analysis of the errors in one image of the set showed that this was the case for 25% of the errors made. 15% of the errors could be attributed to the graph partitioner. The remaining errors would require a more exact nuclei model to be eliminated.

4 Discussion and Future Work

In this work we have presented a novel algorithm to split clustered nuclei by combining graph partitioning with a simple object model. We have shown the performance of our method on two sets of synthetic cell nuclei images. Our method is very fast and performs qualitatively comparable or better to other cell nuclei segmentation techniques.

Future work is necessary for a better automation. The required user input should be reduced, e.g. by unsupervised learning of the model parameters. This would allow the model to include more features, which could further improve the quality. Another approach to reducing the user input, but also to allow more sophisticated binarization methods, would be to find a unifying model which can explain all used parameters. For an example, the parameters σ_{grad} and σ_{bin}, used in the computation of the edge weights, should be deduced from the parameters of the automatically computed image model. We tried to find such a relationship, but it did not work in our experiments.

The anisotropic resolution was one of the biggest challenges for this work. The ratio between spacing in axial direction to the lateral directions was as high as 5 in one of our datasets. This negatively affected several aspects of our method. For example, the approximation of sphericity is hard to obtain, as a nucleus extends only about 4 slices in the axial direction. This can make it difficult to distinguish between a sphere, and a half-sphere. One solution for this

issue would be to include more features in the nucleus model, for example the extent of the object in z-direction.

The anisotropic resolution is especially problematic for the gradient-based edge weighting method. Here, the edge weights are a measure for the probability of a change in intensity. However, such a change has a higher probability when moving along the z-axis, because pixels along this axis have a higher spacing. As a result, the weights for the respective edges are too small, favouring cuts of these edges. One idea to cope with this problem could be to adjust the parameter σ_{grad} accordingly. For example, one could use separate parameters depending on the edge orientation.

This work has dealt with image segmentation of 3D images with anisotropic spacing. Another way to consider the images is as a stack of 2D images. This perception relates closely to the image formation process, as the image sensor which digitizes the real image is a 2D array of capacitors. One approach could be to segment all slices separately and then combine the detected 2D objects. This also has the advantage of an easy parallelization. Another idea is to separate the slices for the sphericity measure in the nucleus model: instead of using a three-dimensional sphericity, we could compute a two-dimensional roundness measure for each involved slice.

From a computer scientist's point of view, the problem in the development of image segmentation algorithms is that the problem is ill-posed [8, 28]. For a well-posed problem, a solution exists, is unique, and is stable, i.e. it depends continuously on the data [18]. The image segmentation problem does not have a unique solution, because the digitization process is not injective: multiple different images can be mapped to the same digital image. This is especially the case for the segmentation of live cell nuclei in 3D fluorescence microscopy. The low imaging quality and the anisotropic resolution lead to ambiguities. Even human experts can differ on their choice of the detection and/or the boundary of an object [7, 8].

It is difficult to design a segmentation algorithm when the correct solution is not defined. This problem results from the lack of ground truth for large-scale image datasets. How should the quality of an algorithm be evaluated? How can the quality be optimized? How can it be compared to other algorithms? Some authors resort to subjective quality comparison [45]. Most either compare against hand-segmented datasets or synthetic benchmark datasets [46]. The latter choice is also the path taken in this work. We have optimized our algorithm and the parameters with two sets of synthetic images. While our method shows superior results on these instances, its performance on real-world images has yet to be evaluated.

References

[1] Yousef Al-Kofahi, Wiem Lassoued, William Lee, and Badrinath Roysam. Improved automatic detection and segmentation of cell nuclei in histopathology images. *IEEE Transactions on Biomedical Engineering*, 57(4):841–852, 2010.

[2] Charles-Edmond Bichot and Patrick Siarry. *Graph partitioning.* John Wiley & Sons, 2013.

[3] Yuri Boykov and Gareth Funka-Lea. Graph cuts and efficient N-D image segmentation. *International Journal of Computer Vision*, 70(2):109–131, 2006.

[4] Yuri Boykov and Vladimir Kolmogorov. Computing geodesics and minimal surfaces via graph cuts. In *9th IEEE International Conference on Computer Vision (ICCV 2003)*, pages 26–33. IEEE Computer Society, 2003.

[5] Yuri Boykov and Vladimir Kolmogorov. An experimental comparison of min-cut/max-flow algorithms for energy minimization in vision. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 26(9):1124–1137, 2004.

[6] Yuri Boykov, Olga Veksler, and Ramin Zabih. Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11): 1222–1239, 2001.

[7] Luís Pedro Coelho, Aabid Shariff, and Robert F. Murphy. Nuclear segmentation in microscope cell images: A hand-segmented dataset and comparison of algorithms. In *Proceedings of the 2009 IEEE International Symposium on Biomedical Imaging: From Nano to Macro*, pages 518–521. IEEE, 2009.

[8] Ondřej Daněk. *Graph Cut Based Image Segmentation in Fluorescence Microscopy.* PhD thesis, Masaryk University, 2012.

[9] Ondřej Daněk and Pavel Matula. On euclidean metric approximation via graph cuts. In Paul Richard and José Braz, editors, *Computer Vision, Imaging and Computer Graphics. Theory and Applications*, volume 229 of *Communications in Computer and Information Science*, pages 125–134. Springer, 2010.

[10] Andrew Delong and Yuri Boykov. A scalable graph-cut algorithm for N-D grids. In *2008 IEEE Computer Society Conference on Computer Vision and Pattern Recognition.* IEEE Computer Society, 2008.

[11] Arthur P. Dempster, Nan M. Laird, and Donald B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–38, 1977.

[12] Ralf Diekmann, Robert Preis, Frank Schlimbach, and Chris Walshaw. Shape-optimized mesh partitioning and load balancing for parallel adaptive FEM. *Parallel Computing*, 26 (12):1555–1581, 2000.

[13] Charles M. Fiduccia and Robert M. Mattheyses. A linear-time heuristic for improving network partitions. In *19th Design Automation Conference*, pages 175–181. ACM/IEEE, 1982.

[14] Michael R. Garey, David S. Johnson, and Larry J. Stockmeyer. Some simplified np-complete problems. In Robert L. Constable, Robert W. Ritchie, Jack W. Carlyle, and Michael A. Harrison, editors, *Proceedings of the 6th Annual ACM Symposium on Theory of Computing*, pages 47–63. ACM, 1974.

[15] Andrew V. Goldberg, Sagi Hed, Haim Kaplan, Robert Endre Tarjan, and Renato Fonseca F. Werneck. Maximum flows by incremental breadth-first search. In Camil Demetrescu and Magnús M. Halldórsson, editors, *Proceedings of the 19th Annual European Symposium on Algorithms*, volume 6942 of *Lecture Notes in Computer Science*, pages 457–468. Springer, 2011.

[16] Rafael C. Gonzalez and Richard E. Woods. Image processing. *Digital image processing*, 2, 2007.

[17] Dorothy M. Greig, Bruce T. Porteous, and Allan H. Seheult. Exact maximum a posteriori estimation for binary images. *Journal of the Royal Statistical Society. Series B (Methodological)*, 51(2):271–279, 1989.

[18] Jacques Hadamard. Sur les problèmes aux dérivées partielles et leur signification physique. *Princeton university bulletin*, 13(49-52):28, 1902.

[19] Yong He, Hui Gong, Benyi Xiong, Xiaofeng Xu, Anan Li, Tao Jiang, Qingtao Sun, Simin Wang, Qingming Luo, and Shangbin Chen. icut: an integrative cut algorithm enables accurate segmentation of touching cells. *Scientific Reports*, 5, 2015.

[20] Bruce Hendrickson and Robert W. Leland. A multi-level algorithm for partitioning graphs. In Sidney Karin, editor, *Supercomputing '95*, page 28. IEEE Computer Society / ACM, 1995.

[21] Jan Huisken, Jim Swoger, Filippo Del Bene, Joachim Wittbrodt, and Ernst HK Stelzer. Optical sectioning deep inside live embryos by selective plane illumination microscopy. *Science*, 305(5686):1007–1009, 2004.

[22] Laurent Hyafil and Ronald L. Rivest. Graph partitioning and constructing optimal decision trees are polynomial complete problems. Technical report, IRIA – Laboratoire de Recherche en Informatique et Automatique, 1973.

[23] L. Ibanez, W. Schroeder, L. Ng, and J. Cates. *The ITK Software Guide*. Kitware, Inc., second edition, 2005. ISBN 1-930934-15-7.

[24] Shinya Inoué. Foundations of confocal scanned imaging in light microscopy. In *Handbook of biological confocal microscopy*, pages 1–19. Springer, 2006.

[25] Ondrej Jamriska, Daniel Sýkora, and Alexander Hornung. Cache-efficient graph cuts on structured grids. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 3673–3680. IEEE Computer Society, IEEE Computer Society, 2012.

[26] George Karypis and Vipin Kumar. A fast and high quality multilevel scheme for partitioning irregular graphs. *SIAM Journal on scientific Computing*, 20(1):359–392, 1998.

[27] Philipp J. Keller, Annette D. Schmidt, Joachim Wittbrodt, and Ernst H.K. Stelzer. Reconstruction of zebrafish early embryonic development by scanned light sheet microscopy. *Science*, 322(5904):1065–1069, 2008. doi: 10.1126/science.1162493.

[28] Khaled Khairy and Philipp J. Keller. Reconstructing embryonic development. *Genesis*, 49(7):488–513, 2011.

[29] Gang Lin, Umesh Adiga, Kathy Olson, John F. Guzowski, Carol A. Barnes, and Badrinath Roysam. A hybrid 3d watershed algorithm incorporating gradient cues and object models for automatic segmentation of nuclei in confocal image stacks. *Cytometry Part A*, 56A(1): 23–36, 2003.

[30] Gang Lin, Monica K. Chawla, Kathy Olson, John F. Guzowski, Carol A. Barnes, and Badrinath Roysam. Hierarchical, model-based merging of multiple fragments for improved three-dimensional segmentation of nuclei. *Cytometry Part A*, 63(1):20–33, 2005.

[31] Jiangyu Liu and Jian Sun. Parallel graph-cuts by adaptive bottom-up merging. In *The Twenty-Third IEEE Conference on Computer Vision and Pattern Recognition*, pages 2181–2188. IEEE Computer Society, 2010.

[32] Lifeng Liu and Stan Sclaroff. Region segmentation via deformable model-guided split and merge. In *ICCV*, pages 98–104. IEEE Computer Society, 2001.

[33] Xinghua Lou, Ullrich Köthe, Jochen Wittbrodt, and Fred A. Hamprecht. Learning to segment dense cell nuclei with shape prior. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, pages 1012–1018. IEEE Computer Society, IEEE Computer Society, 2012.

[34] Marvin Minsky. Microscopy apparatus, December 19 1961. US Patent 3,013,467.

[35] Michiel Muller. *Introduction to confocal fluorescence microscopy*, volume 69. SPIE press, 2006.

[36] Nobuyuki Otsu. A threshold selection method from gray-level histograms. *Automatica*, 11 (285-296):23–27, 1975.

[37] Christophe Restif. *Segmentation and evaluation of fluorescence microscopy images*. PhD thesis, Oxford Brookes University, 2006.

[38] Christophe Restif. Towards safer, faster prenatal genetic tests: Novel unsupervised, automatic and robust methods of segmentation of nuclei and probes. In Ales Leonardis, Horst Bischof, and Axel Pinz, editors, *9th European Conference on Computer Vision, 2006*, volume 3954 of *Lecture Notes in Computer Science*, pages 437–450. Springer, 2006.

[39] T.W. Ridler and S. Calvard. Picture thresholding using an iterative selection method. *IEEE transactions on Systems, Man and Cybernetics*, 8(8):630–632, 1978.

[40] Peter Sanders and Christian Schulz. Engineering multilevel graph partitioning algorithms. In Camil Demetrescu and Magnús M. Halldórsson, editors, *Proceedings of the 19th Annual European Symposium on Algorithms*, volume 6942 of *Lecture Notes in Computer Science*, pages 469–480. Springer, 2011.

[41] Peter Sanders and Christian Schulz. Think locally, act globally: Highly balanced graph partitioning. In *Proceedings of the 12th International Symposium on Experimental Algorithms*, volume 7933 of *Lecture Notes in Computer Science*, pages 164–175. Springer, 2013.

[42] Jianbo Shi and Jitendra Malik. Normalized cuts and image segmentation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 22(8):888–905, 2000.

[43] Milan Sonka, Vaclav Hlavac, and Roger Boyle. *Image processing, analysis, and machine vision*. Cengage Learning, 2014.

[44] Johannes Stegmaier, Arif ul Maula Khan, Markus Reischl, and Ralf Mikut. Challenges of uncertainty propagation in image analysis. In *Proceedings of the 22nd Workshop on Computational Intelligence*, pages 55–69. KIT Scientific Publishing, 2012.

[45] Johannes Stegmaier, Jens C. Otte, Andrei Kobitski, Andreas Bartschat, Ariel Garcia, G. Ulrich Nienhaus, Uwe Strähle, and Ralf Mikut. Fast segmentation of stained nuclei in terabyte-scale, time resolved 3d microscopy image stacks. *PloS one*, 9(2):e90036, 2014.

[46] Johannes Stegmaier, Julian Arz, Benjamin Schott, Jens C. Otte, Andrei Kobitski, G. Ulrich Nienhaus, Uwe Strähle, Peter Sanders, and Ralf Mikut. Generating semi-synthetic validation benchmarks for embryomics. In *Proceedings of the 13th IEEE International Symposium on Biomedical Imaging*, pages 684–688. IEEE, 2016.

[47] David Svoboda, Michal Kozubek, and Stanislav Stejskal. Generation of digital phantoms of cell nuclei and simulation of image formation in 3d image cytometry. *Cytometry Part A*, 75(6):494–509, 2009.

[48] Richard Szeliski. *Computer vision: algorithms and applications*. Springer Science & Business Media, 2010.

[49] Raju Tomer, Khaled Khairy, Fernando Amat, and Philipp J. Keller. Quantitative high-speed imaging of entire developing embryos with simultaneous multiview light-sheet microscopy. *Nature Methods*, 9(7):755–763, 2012.

[50] Jayaram K. Udupa and Supun Samarasekera. Fuzzy connectedness and object definition: Theory, algorithms, and applications in image segmentation. *CVGIP: Graphical Model and Image Processing*, 58(3):246–261, 1996.

[51] Hakon Wadell. Volume, shape, and roundness of quartz particles. *The Journal of Geology*, pages 250–280, 1935.